

07-03-00

jc714 U.S. PTO

09/608395

jc714 U.S. PTO
06/30/00

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventorship..... Larson et al.
 Applicant..... Microsoft Corporation
 Attorney's Docket No. MS1-479US
 Title: Partial Pre-Aggregation In Relational Database Queries

TRANSMITTAL LETTER AND CERTIFICATE OF MAILING

To: Commissioner of Patents and Trademarks
 Washington, D.C. 20231

From: James R. Banowsky (509) 324-9256
 Lee & Hayes, PLLC
 421 W. Riverside Avenue, Suite 500
 Spokane, WA 99201

The following enumerated items accompany this transmittal letter and are being submitted for the matter identified in the above caption.

1. Transmittal Letter with Certificate of Mailing included.
2. PTO Return Postcard Receipt
3. Check in the Amount of \$ 1024.00
4. Fee Transmittal
5. New patent application (title page plus 35 pages, including claims 1-26 & Abstract)
6. Executed Declaration
7. 8 sheets of formal drawings (Figs. 1-10)
8. Assignment w/Recordation Cover Sheet

Large Entity Status [x]

Small Entity Status []

The Commissioner is hereby authorized to charge payment of fees or credit overpayments to Deposit Account No. 12-0769 in connection with any patent application filing fees under 37 CFR 1.16, and any processing fees under 37 CFR 1.17.

Date: 6/30/00

By:

James R. Banowsky
 James R. Banowsky
 Reg. No. 37,773

CERTIFICATE OF MAILING

I hereby certify that the items listed above as enclosed are being deposited with the U.S. Postal Service as either first class mail, or Express Mail if the blank for Express Mail No. is completed below, in an envelope addressed to The Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the below-indicated date. Any Express Mail No. has also been marked on the listed items.

Express Mail No. (if applicable) EL624352927

Date: 6/30/2000

By:

Dana L. Calhoun
 Dana L. Calhoun

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Partial Pre-Aggregation In Relational Database Queries

Inventor(s):

Per-Åke Larson
César A. Galindo-Legaria

ATTORNEY'S DOCKET NO. MS1-479US

TECHNICAL FIELD

This invention relates to relational database systems and, more particularly, to relational database queries utilizing aggregation operations.

BACKGROUND

Relational database systems are a type of database or database management system that stores information in tables -- rows and columns of data. Typically, the rows of a table represent records (collections of information about separate items) and the columns represent fields (particular attributes of a record). An example of a relational database system is the SQL (Structured Query Language) Server database system manufactured and sold by Microsoft Corporation.

Database programs utilize queries to perform searches on one or more databases. Queries are composed of operators that perform a function involving one or more tables. One particular type of query that is frequently used in relational database systems groups records according to the value of one or more columns in the records. A query of this type is often referred to as an aggregation operation, an aggregation query or, simply, an aggregation ("GROUP BY" in SQL). As an example, suppose a business administrator wants to compute a sales total for a set of customers from a number of invoices. Each record represents one invoice and contains, among other things, a customer number and a dollar amount. The administrator might formulate an aggregation query that groups the invoices according to customer number and sums the dollar amounts on the invoices for each customer.

Other operators can be utilized together with an aggregation to perform a more sophisticated query on one or more tables. One such operator is a join

operation, also referred to as a join query or a join. A join operation takes information in one table and combines that information with related information in another table. In the example given above, suppose the administrator would also like to see the name of the customer and the phone number of a customer contact (person) for each customer together with the total sales information for the customer. Suppose, also, that the customer name and customer contact information are not listed in the sales table that contains the sales figures - they are listed in a customer table, which contains, among other things, the customer name and contact information for the customer.

In this case, a join operation is first performed to combine the customer name and contact information with the invoice information. Subsequently, the aggregation operation calculates the total sales per customer as outlined above, and the administrator has a result that shows the customer name, the phone number of the customer contact person, and the total sales for that customer. There is, however, a significant cost to performing the join followed by the aggregation, due to the time necessary to perform certain input/output (I/O) operations.

The data records of a database are commonly stored on disk arrays or other forms of non-volatile memory. Queries performed on relational databases, such as the aggregation query, require that all data records be loaded into volatile memory (*i.e.*, random access memory, or "RAM") for processing. However, relational databases often contain a large amount of data, which surpasses the volatile memory resources. As a result, records are loaded into volatile memory in batches to create a "record store," and large intermediate results must frequently be written to non-volatile memory such as a disk. Transferring large amounts of data

1 between volatile and non-volatile memory significantly increases the cost of
2 processing a query.

3 This cost problem is compounded in the case of the query having a join
4 followed by an aggregation. Consider the situation if the sales table contains one
5 million invoices and there are ten thousand different customer records in the
6 customer table. It is doubtful that there is sufficient RAM to load both tables
7 completely into RAM and perform the join in RAM. Therefore, joining the
8 customer table with the sales table requires loading a record from the sales table,
9 identifying the customer number in that record, searching the customer table to
10 locate a record for that customer, loading the record for that customer into RAM,
11 and creating a new record combining the sales and customer information. This
12 requires the computer to access the disk for each record in the sales database, or
13 one million times. The join algorithm outlined above is known as a (simple)
14 nested-loop join algorithm. It is relatively inefficient but it is used here to outline
15 the basic idea of join processing. There are other more efficient join algorithms
16 (hash join, merge join, etc.) but they are considerably more complex and will not
17 be discussed in detail herein. Those skilled in the art will be familiar with the
18 intricacies of these algorithms.

19 The bulk of the processing overhead in the query described above is related
20 to the join operation. This is clear, since the join operation requires accessing the
21 disk for each record. In this example, the disk is accessed one million times. The
22 time required for a disk access operation is significantly large in comparison to
23 other computer operations. Therefore, the join operation and the disk accesses
24 required therein are prime targets for any endeavor to reduce processing overhead.
25

1 If the records used for input to the join operation are reduced, then the
2 overhead of the join operation is also reduced. One way in which the input
3 records to a join operation can be reduced is to perform an additional aggregation -
4 a "pre-aggregation" - prior to computing the join. Continuing on the example
5 query, the total sales per customer could be computed first by processing only the
6 sales table, then joining the resulting table with the customer table.

7 While this reduces the input to the join, significant overhead is still required
8 for the pre-aggregation. This is because a pre-aggregation operation is held to the
9 same standard as any aggregation operation. That is, the pre-aggregation
10 operation performs a complete aggregation, producing just a single output record
11 for each customer.

12 13 SUMMARY

14 The implementations described herein concern a "partial pre-aggregation"
15 operation that is similar to an aggregation, but that does not necessarily continue
16 the aggregation process until there is only one output record for each group. The
17 partial pre-aggregation provides a result that reduces the number of records for
18 input into a subsequent operation, but requires less overhead than a complete pre-
19 aggregation.

20 Referring to the example above, if partial pre-aggregation is performed on
21 the records prior to performing the join, the number of records input to the join
22 operation can be significantly reduced. If the partial pre-aggregation happens to
23 output an average of five records for every customer, the records input to the join
24 will be reduced from one million to fifty thousand. This is not as complete as a
25

1 full pre-aggregation, which would reduce the number of records to ten thousand.
2 However, the cost savings are significant.

3 In a partial pre-aggregation, when a new input record (representing an
4 invoice) is received from the disk into RAM, a determination is made as to
5 whether the input record belongs to the same customer as an aggregation record
6 already in a record store created for the partial pre-aggregation. (An aggregation
7 record stores the customer number and a running total of the sales for that
8 customer.) If so, the new record is combined with the aggregation record in the
9 record store. Therefore, what were once two records is now one. If the input
10 record does not belong to the same customer as any of the aggregation records
11 already in the record store, a new aggregation record for that customer is created
12 in the record store, if there is sufficient space for another record in the record
13 store. If the record store is full, space is vacated by outputting one or more of the
14 aggregation records to a subsequent operation. Note that aggregation records are
15 never output to disk as part of the partial pre-aggregation process but are
16 immediately passed on to the next operator.

17 It is not always most efficient to perform a partial pre-aggregation on a
18 record store prior to a subsequent operation, such as a join. Therefore, one
19 implementation described herein includes a query optimizer that provides an
20 estimate of the number of records that will be output from a partial pre-
21 aggregation. If the query optimizer suggests that the cost of a partial pre-
22 aggregation will outweigh the benefits, the query will process the records without
23 performing a partial pre-aggregation. If the query optimizer suggests that a partial
24 pre-aggregation will be more efficient, then the partial pre-aggregation is
25 performed as a part of the query.

1 In another implementation, a special case is described in which the
2 grouping columns are the same for the partial pre-aggregation and the join. In
3 such a case, the partial pre-aggregation and the join can be processed together, as
4 opposed to them being processed sequentially.

5 In this implementation, a batch of records is input to a record store up to a
6 capacity of records that the record store can store (or until there are no more
7 records to input). If the query optimizer determines that a partial pre-aggregation
8 is in order, aggregation is performed on the records as they are loaded into the
9 record store. This reduces the number of records in the record store to one record
10 for each customer represented by a record in the record store. (Note that a
11 customer represented by a single record in the record store may also be
12 represented by several more records waiting to be input into the record store; thus,
13 this is a partial pre-aggregation.)

14 Once the pre-aggregation operator has filled the record store with
15 aggregation records, a join is performed using the records in the record store and
16 the resulting joined records are output to a subsequent operator. This prevents the
17 join operation from having to create its own record store, thereby saving a
18 significant amount of memory. When the join is completed, the records in the
19 record store are discarded. If the subsequent operator is not an aggregation, then
20 an aggregation will ultimately be performed on all the records to produce the
21 desired output.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows a relational database computer system.

Fig. 2 is a flow diagram of a method utilizing a partial pre-aggregation operation.

Fig. 3 is a flow diagram of a partial pre-aggregation operation.

Fig. 4 is an illustration of a sales table.

Fig. 5 is an illustration of a customer table.

Fig. 6a is an illustration showing a first group of aggregation records output to a join operator as a result of applying partial pre-aggregation to the sales table.

Fig. 6b is an illustration showing a second group of aggregation records output to a join operator as a result of applying partial pre-aggregation to the sales table.

Fig. 7 is an illustration showing a result of applying a join operation to the result of the partial pre-aggregation and the customer table.

Fig. 8 is an illustration showing a result of aggregating the result of the join operation.

Fig. 9 is a flow diagram of a method utilizing a single record store for a partial pre-aggregation operation and a join operation.

Fig. 10a is an illustration showing a first group of records input to a record store.

Fig. 10b is an illustration showing records resulting from performing a partial pre-aggregation on the record store.

Fig. 10c is an illustration showing records resulting from performing a join on the aggregated record store.

1 Fig. 11a is an illustration showing a second group of records input to a
2 record store.

3 Fig. 11b is an illustration showing records resulting from performing a
4 partial pre-aggregation on the record store.

5 Fig. 11c is an illustration showing records resulting from performing a join
6 on the aggregated record store.

7 Fig. 12 is an illustration showing records resulting from performing an
8 aggregation on the records output from the join operations.

9 10 **DETAILED DESCRIPTION**

11 Fig. 1 shows a relational database system 100 having a computer 102 and a
12 non-volatile memory 104 interfaced with the computer 102. The computer 102
13 has a processing unit 106 and a main memory 108. The main memory 108 is
14 volatile memory and can be implemented, for example, as volatile RAM (Random
15 Access Memory). The non-volatile memory 104 provides permanent storage for
16 relational database records. The non-volatile memory 104 can be implemented in
17 a variety of ways, including disk arrays, disk drives (e.g., hard and floppy),
18 read/write CD ROMS, tape backups, reel-to-reel, and the like.

19 The relational database system 100 is shown in an operational state in
20 which a relational database program 110 is loaded in main memory 108 for
21 execution on the processing unit 106. The relational database program 110 is
22 permanently stored on non-volatile memory 104 and loaded into the main memory
23 108 when launched. An example of a relational database program is the SQL
24 Server program sold by Microsoft Corporation. It is also noted that aspects of this
25 invention concerning query processing may be used in other types of programs

1 that may employ relational database concepts, such as spreadsheet programs,
2 accounting software, workflow management software, and the like.

3 The relational database program 110 has a query processor 112, which is a
4 program that implements aspects of the embodiments described herein. The
5 relational database program also includes a query optimizer 114, which examines a
6 query and the data to be processed by the query to determine the optimum manner
7 in which the query should be processed.

8 Some operators, in particular, join and aggregation operators, require main
9 memory space for storing records during processing. This type of working storage
10 is here referred to as a "record store". A portion of the main memory 108 is
11 shown as being assigned to record stores 116 created and used by various operators.
12 Note that records stores are not permanent but created and destroyed by operators
13 as needed. Neither are they all of the same size – different operators may use
14 record stores of different size. The relational database system 110 is representative
15 of many diverse implementations, including a stand-alone computer, a database
16 server for a network of PCs or workstations, an online server for Internet service
17 providers, a mainframe computing system, and the like. The relational database
18 system 110 runs on top of an operating system (not shown), which is preferably a
19 multitasking operating system that allows simultaneous execution of multiple
20 applications or multiple threads of one or more applications. Examples of suitable
21 operating systems include a Windows brand operating system sold by Microsoft
22 Corporation, such as the Windows NT workstation operating system, as well as
23 UNIX based operating systems.

24 Fig. 2 is a flow diagram of a method utilizing partial pre-aggregation
25 according to the embodiments described herein. At step 200, the query processor

1 112 of the relational database program 110 begins to process a query. For
2 discussion purposes, the query includes a join operation followed by an
3 aggregation operation.

4 The relational database program 110 invokes the query optimizer 114 to
5 examine the query and the data to be processed by the query to determine if the
6 query is one on which partial pre-aggregation can be performed (step 202). If the
7 query includes an aggregation operation, then it may be possible to perform a
8 partial pre-aggregation as part of the query. Partial pre-aggregation can be applied
9 wherever there is a column set that functionally determines the final grouping
10 columns. In other words, whenever grouping and aggregating is done on a column
11 set "G," pre-aggregation can be performed on any column set that functionally
12 determines "G." Determining when a pre-aggregation can be performed is known
13 in the art.

14 If it is not possible to perform a partial pre-aggregation, then there is no
15 need to determine if partial pre-aggregation should be done to economize
16 overhead. In that case, the query processor 112 simply continues to process the
17 query without performing a partial pre-aggregation at step 208 ("NO" branch, step
18 202). If, however, it is possible to perform a partial pre-aggregation ("YES"
19 branch, step 202), the query optimizer 114 analyzes the query to estimate the
20 benefits of running a partial pre-aggregation as opposed to the costs of performing
21 a partial pre-aggregation (step 204).

22 Several mathematical models are available and known in the art that can be
23 used for this estimation. However, a refined mathematical model is described
24 below that may be used to provide more accurate estimates than previous models.
25

Mathematical Model

If an input stream (records that are input to be processed) contains D distinct groups, complete aggregation reduces the input to exactly D output records. Partial aggregation will output more than D records. Exactly how many more depends on several factors, such as the amount of memory available, the number of groups, the group size distribution, and the ordering of the input. For the present example, it is assumed that the input is a stream of randomly selected records.

Assume that the input records are divided among D different groups, labeled $1, 2, \dots, D$, and that p_i denotes the probability that a record belongs to group i . (p_1, p_2, \dots, p_D , is a group size distribution). The actual group labels (grouping column values) do not matter for mathematical model purposes, but it is assumed that $p_1 \geq p_2 \geq p_D$.

Group size distribution is modeled as a generalized Zipf distribution that is defined by:

$$p_i = 1/c)(1/i)^\alpha, i=1,2,\dots,D \qquad c=\sum_{i=1}^D (1/i)^\alpha$$

where α is a positive constant.

Setting $\alpha = 1$ gives the traditional Zipf distribution and $\alpha = 0$ gives a uniform distribution. Increasing α increases the skew in the distribution, which increases the data reduction obtained by partial aggregation.

An input record will either be absorbed by a group already in memory or will create a new group. Group labels are modeled as being independently and

randomly drawn from the distribution p_1, p_2, \dots, p_D . The expected number of distinct group labels occurring in a sample of n records (where n denotes the number of records read so far) equals:

$$G(n) = D - \sum_{i=1}^D (1 - p_i)^n$$

where $(1 - p_i)^n$ is the probability that no record with group label i occurs among the n input records. Note that the function G is also well defined for non-integer arguments.

An absorption rate at point n - the probability that record $n+1$ will be absorbed into one of the groups already in memory - is calculated as:

$$A(n) = 1 - (G(n+1) - G(n)) = \sum_{i=1}^D (1 - p_i)^n$$

where $G(n+1) - G(n)$ is the probability that record $n+1$ will *not* be absorbed.

The number of input records expected to be processed before the memory reaches capacity ($R(M)$) (assuming that there is memory space for storing, at most, M group records and that $M < D$) is the inverse of the function G :

$$R(M) = G^{-1}(M).$$

$R(M)$ is computed by solving $M = G(X)$ for X .

$R(M)$ is substituted into function A , below, to obtain an estimate of the absorption rate obtained when storing M group records in memory. This function equals:

$$A(R(M)) = 1 - \sum_{i=1}^D (1 - p_i)^{R(M)}.$$

$A(R(M))$ is a measure of the “absorption power” of memory space for M group records. If the input consists of N records ($N > D$), the number of output records can be estimated as:

$$T(N) = M + (N - M)(1 - A(R(M))) = M + (N - M) \sum_{i=1}^D (1 - p_i)^{R(M)}.$$

This function applies to replacement policies that always leave the last M distinct values encountered in the input stream in memory.

$T(N)$ provides the number of output records for N input records. Therefore, an estimate can be made of the number of output records when the number of input records is known or can be reasonably estimated. If the estimated number of output records is significantly less than the number of input records, it is beneficial to apply partial pre-aggregation.

Note that for the special case of uniform distribution ($p_i = 1/D$), closed formulas for the four functions are:

$$G_U(n) = D(1 - 1/D)^n$$

$$R_U(M) = \log_{(1 - 1/D)}(1 - M/D) = \ln(1 - M/D)/\ln(1 - 1/D) \text{ where } M = G^{-1}.$$

$$A_U(R_U(M)) = M/D$$

$$T_U(N) = M + (N - M)(1 - M/D)$$

Referring now back to Fig. 2, if the query optimizer 114 indicates that it would be cost efficient to perform a partial pre-aggregation (“YES” branch, step 204), then the query processor 112 performs the partial pre-aggregation at step 206. Details of the partial pre-aggregation will be discussed below with reference to Fig. 3. If the query optimizer 114 determines that no cost savings will be

1 realized from performing a partial pre-aggregation (“NO” branch, step 204), then
2 the query processing is continued at step 208.

3 It is noted that the partial pre-aggregation operator is a “non-blocking” or
4 “streaming” operator, which means that it may generate output records before it
5 receives all the input records it is to process. This is a significant advantage of
6 partial pre-aggregation, since it can run concurrently with a subsequent operator,
7 such as a join. Therefore, it is noted that step 206 - “Perform Partial Pre-
8 Aggregation” - may be processed concurrently as a part of the query and not
9 necessarily sequentially, or separate from the processing of other operators. In
10 other words, other operators may be running concurrently with the partial pre-
11 aggregation and, therefore, would not necessarily fall sequentially within the flow
12 chart of Fig. 2.

13 The steps depicted in Fig. 2 will be revisited below, with reference to a
14 specific example outlined in Figs. 4-8.

15 Fig. 3 is a flow diagram depicting the steps of a partial pre-aggregation.
16 The steps outlined in Fig. 3 comprise the details of step 206 of Fig. 2.

17 At step 300 of Fig. 3, the relational database program 110 performs a check
18 to determine if there are input records available for processing. If so (“YES”
19 branch, step 300), an input record is received by the partial pre-aggregation
20 operator at step 302. This input record may be stored in a partitioned area of the
21 main memory 108, depending on the implementation used. Other memory
22 partitions may be utilized for processing programs, hashing, partitioning, etc.

23 If the value of the grouping column(s) of the input record matches the
24 grouping column(s) of an aggregation record already in the operators record store
25

1 116 (step 304), then the input record is combined with the matching record at step
2 306 and the process reverts to step 300.

3 If, however, no match is found ("NO" branch, step 304), then a new
4 aggregation record matching the input record must be added to the record store
5 116. Before that can be done, the relational database program 110 determines if
6 the record store 116 is full at step 308, the record store 116 having a capacity to
7 store a limited number of records. It will be appreciated that the record store 116
8 may actually hold any practicable number of records.

9 If the record store 116 is not full ("NO" branch, step 308), then the new
10 aggregation record is created and stored in the record store 116 at step 312 and the
11 process again reverts to step 300. If the record store 116 is full ("YES" branch,
12 step 308), then some of the existing aggregation records in the record store 116 are
13 output to the subsequent join operator at step 310.

14 Note that the record store 116, at this point, contains only one record for
15 *each value* of the grouping column (customer number). Once the pre-aggregation
16 operator has output an aggregation record to the join operator, the record is deleted
17 from the record store. Suppose the record output was related to a customer with
18 customer number 1234. The input stream may contain additional invoices for
19 customer 1234, in which case a new aggregation record will be created for
20 customer 1234 and eventually output to the join operator. In other words, the
21 output stream from pre-aggregation may contain multiple records related to the
22 same customer, each one covering a subset of that customer's invoices.
23 Traditional, complete aggregation always outputs a single record for each
24 customer. This is the difference between *partial* pre-aggregation and pre-
25 aggregation. However, this will not affect the final output since a final

1 aggregation will ultimately be performed on the results of the partial pre-
2 aggregation.

3 When there are no more records to be input ("NO" branch, step 300), all
4 remaining aggregation records in the record store are output to the join operator.
5 This ensures that all the records are eventually output to the join operator.

6 The theory of partial pre-aggregation can be more easily explained using a
7 practical example. Suppose that a business administrator wants an output that lists
8 a total of all the sales for a certain year made to each customer of the business,
9 together with a contact (person) for each customer and a telephone number for the
10 contact.

11 The administrator has a database system that includes a sales table for the
12 year desired that contains a record for each invoice, each record showing, among
13 other things, the customer number and the amount of sales for that particular
14 invoice. The database system also includes a customer table that contains a record
15 for each customer of the business, each record showing, among other things, the
16 contact for that customer and a telephone number for that contact. Note that since
17 the administrator needs information from two tables, there will be a join operation
18 in the query. Also, since the information sought by the administrator includes
19 grouping by customer number, there will be an aggregation operation in the query.

20 Fig. 4 shows a sales table 400 for the example outlined above. The sales
21 table 400 includes six rows, or records, identified as record 1 402, record 2 404,
22 record 3 406, record 4 408, record 5 410 and record 6 412. It is noted that
23 although six records are shown, the sales database can contain any practicable
24 number of records.
25

1 The sales table 400 also includes four columns, or fields. These are
2 customer number 414, invoice date 416, due date 418 and invoice amount 420.
3 (The invoice date 416 and due date 418 are not shown as they are not relevant to
4 the following discussion).

5 For record 1 402, the customer number 414 is "9810" and the invoice
6 amount 420 is \$100. For record 2 404, the customer number 414 is "9815" and
7 the invoice amount 420 is \$200. For record 3 406, the customer number 414 is
8 "9810" and the invoice amount 420 is \$300. For record 4 408, the customer
9 number 414 is "9821" and the invoice amount 420 is \$400. For record 5 410, the
10 customer number 414 is "9810" and the invoice amount 420 is \$500. For record 6
11 412, the customer number 414 is "9821" and the invoice amount 420 is \$600.

12 Fig. 5 shows a customer table 500 for the example outlined above. The
13 customer table 500 includes four rows, or records, identified as record 1 502,
14 record 2 504, record 3 506, and record 4 508. It is noted that although six records
15 are shown, the sales table can contain any practicable number of records.

16 The customer table 500 also includes four columns, or fields. These are
17 customer number 510, customer name 512, customer contact 514, and contact
18 phone 516. (The customer names 514 are not shown because they are not relevant
19 to the following discussion). For record 1 502, the customer number 510 is
20 "9810," the customer contact 514 is "Adams" and the contact phone is 123-4567.
21 For record 2 504, the customer number 510 is "9815," the customer contact 514 is
22 "Bethard" and the contact phone 516 is 345-6789. For record 3 506, the customer
23 number 510 is "9819," the customer contact 514 is "Gill" and the contact phone
24 516 is 987-6543. For record 4 508, the customer number 510 is "9821," the
25 customer contact 514 is "Danner" and the contact phone 516 is 765-4321.

1 Referring back to Fig. 2 with continuing reference to the sales table 400 of
2 Fig. 4 and the customer database 500 of Fig. 5, the query processing begins at step
3 200. At step 202, the query processor 112 determines that a partial pre-
4 aggregation can be performed since the query requires grouping columns, i.e., an
5 aggregation.

6 At step 204, the query optimizer 114 determines that it would be cost
7 efficient to perform a partial pre-aggregation on the sales table prior to the join
8 (assumed in this example). The partial pre-aggregation is executed at step 206.

9 Fig. 3 shows the details of the partial pre-aggregation, which begins at step
10 300, when the query processor 112 verifies that there are input records to load into
11 the record store 116, created for this purpose by the pre-aggregation operator

12 At step 302, record 1 402 of the sales table 400 is input into the main
13 (volatile) memory 108, specifically, the record store 116. At step 304, the value of
14 the grouping column (customer number 414 "9810" for this example) of record 1
15 402 is checked for a match with an aggregation record already in the record store
16 116. Since no other records have been processed yet, there is no match. The
17 query processor 112 then determines whether the record store 116 is full at step
18 308. For this example, assume that the record store 116 can contain only two
19 records, although it is noted that, in practice, the record store 116 can contain a
20 much larger number of records. For this example, it is more convenient to discuss
21 the record store 116 as having a small capacity.

22 After confirming that there is room in the record store 116 for another
23 record, record 1 402 is added to the record store 116 (step 312) and the process
24 reverts to step 300, where it is determined that there are more input.
25

At step 302, record 2 404 of the sales table 400 is input into record store 116. At step 304, the value of the grouping column of record 2 404 ("9815") is checked for a match with an aggregate record in the record store 116. Since the only record in the record store 116 has the customer number "9810," there is no match. Therefore, record 2 404 is added to the record store 116 at step 312 after determining that there is room in the record store 116 (step 308). The process reverts again to step 300.

There are still input records to process, so at step 302, record 3 406 is input into the record store 116. At step 304, the value of the grouping column of record 3 406 ("9810") is checked for a matches with an aggregate record in the record store 116. The value of the grouping column for record 3 matches the record in the record store 116 having the value "9810" for the grouping column. Therefore, at step 306, the records are combined by summing the invoice amount 420 fields of the records.

More records are available to be input at step 300, so record 4 408 is input into the record store 116 at step 302. The grouping column value of record 4 408 ("9821") does not match a grouping column value of any record in the record store 116 ("9810," "9815"). The record store 116 is checked at step 308 to determine if it has reached capacity. Since there are now two records in the record store 116, the record store 116 cannot accept record 4 408. Therefore, at least one aggregation record in the record store 116 is output to the subsequent join operator at step 310. Assuming all the aggregation records in the record store 116 are output to the join, the record store 116 is now empty. Therefore, at step 312, record 4 408 is added to the record store 116.

Fig. 6a shows the first group of aggregation records output from the record store 116 to the join operator. The aggregation records are record 1 602 and record 2 604. Record 1 602 contains the customer number 606 "9810" and a total invoice amount 608 of "\$400." The figure of \$400 was derived from adding the first record for customer number "9810," which had an invoice amount of \$100, to the second record for customer number "9810," which had an invoice amount of \$200. Record 2 604 contains the customer number 606 "9815" and a total invoice amount 608 of "\$200." This is the same as the only record previously stored for customer number 9815.

Referring back to Fig. 3, step 300, there are more records to input, so at step 302, record 5 410 is received into the record store 116 of the main memory 108. At step 304, the value of the grouping column of record 5 410 ("9810") is checked against the aggregation records in the record store 116 for a matching grouping column value. Since all records having customer number "9810" as the grouping column value have been output to the join operator, no match is found. The record store 116 is checked for capacity at step 308 and, since it is not full, record 5 410 is added to the record store 116 at step 312.

There are more records to input (step 300) so record 6 412 is input into the memory pate 116. At step 304, the value of the grouping column of record 6 412 ("9821") is checked for a match with the value of the grouping column of the aggregation records in the record store 116. Since there is a record in the record store 116 that has the grouping column value "9821," record 6 412 is combined with the aggregation record in the record store 116 and the invoice amounts 420 for record 6 412 and the matching aggregation record are summed into one record.

At step 300, there are no more input records to process, so the aggregation records remaining in the record store 116 are output to the join operator at step 314. In this case, the record store 116 happens to be filled to capacity. However, it is noted that the contents of the record store 116 will be output to the join operator even if it is not full when there are no further records to process.

It is also noted that records may be input into the record store 116 when the record store 116 is full as long as the value of the grouping column of the input record matches the value of the grouping column of an aggregation record in the record store 116. This is because the input record will be combined with the aggregation record and the number of aggregation records will not increase.

Fig. 6b shows the second group of aggregation records output to the join operator. The aggregation records in Fig. 6b are record 1 610 and record 2 612. Record 1 610 contains the customer number 614 "9821" and a total invoice amount 616 of "\$1000." The figure of \$1000 was derived from adding the first record for customer number 9821, which had an invoice amount of \$400, to the second record for customer number 9821, which had an invoice amount of \$600. Record 2 612 contains the customer number 614 ("9810") and the total invoice amount 616 "\$500." This is the same as the only record previously stored in the record store 116 for customer number 9810.

Figs. 6a and 6b have now been output to the join operator. It is clear that the results of the partial pre-aggregation are less complete than the result would be for a complete pre-aggregation or aggregation. This is because there are two records for customer number 9810 that have been output to the join operation.. As will be seen, however, this won't matter in the end because the final aggregation will provide an output conforming to a complete aggregation. It can also be seen

1 in Figs. 6a and 6b that the partial pre-aggregation has significantly reduced the
2 number of records that will be input into the subsequent join operation.

3 If the partial pre-aggregation had not been performed prior to the join
4 operation, there would be six records to input to the join. Performing the partial
5 pre-aggregation reduced that number to four records. This reduction seems small
6 in the example provided, but the reduction is on the order of one-third, which is
7 significant. If there are several thousand or million records to begin with, the
8 significance of the amount of overhead that can be saved by performing a partial
9 pre-aggregation will be appreciated.

10 Now referring back to Fig. 2, at step 208, the query processing continues
11 with the join operation and the aggregation. Fig. 7 shows the join result 700,
12 which has joined the results of the partial pre-aggregation with the customer
13 contact 514 and the contact phone 516 from the customer table 500 (Fig. 5).

14 The join results 700 include four records. Record 1 702 contains customer
15 number 710 "9810," customer contact 712 "Adams," contact phone 714 "123-
16 4567" and total sales 716 "\$400." Record 2 704 contains customer number 710
17 "105," customer contact 712 "Bethard," contact phone 714 "345-6789" and total
18 sales 716 "\$200." Record 3 706 contains customer number 710 "9821," customer
19 contact 712 "Danner," contact phone 714 "987-6543" and total sales 716 "\$1000."
20 Record 4 708 contains customer number 710 "9810," customer contact 712
21 "Adams," contact phone 714 "123-4567" and total sales 714 "\$500."

22 After the join is complete, an aggregation operation is performed on the
23 join results 700. Fig. 8 shows the aggregation result 800, which has three records.
24 Record 1 802 includes customer number 808 "9810," customer contact 810
25 "Adams," contact phone 812 "123-4567" and total sales 814 "\$900." Record 2

1 804 includes customer number 808 "9815," customer contact "Bethard," contact
2 phone 812 "345-6789" and total sales 814 "\$200." Record 3 806 includes
3 customer number 808 "9821," customer contact "Danner," contact phone 812 and
4 total sales 814 "\$1000."

5 The pre-aggregation eliminates a significant number of redundant customer
6 records and the query has provided the output initially requested - the total sales
7 for each customer together with the customer contact and the contact phone
8 number - while drastically reducing the overhead required to process the query.

9 10 Combining Partial Pre-Aggregation and Join

11 For the special case when the grouping columns are the same as the join
12 columns, pre-aggregation can be combined with the join to utilize a single record
13 store. A pre-aggregation is performed on records as they are input to a record
14 store. When the record store is full, the records are immediately joined with other
15 data and the resulting joined records are output to a subsequent operator.

16 Fig. 9 shows that processing a query begins at step 900. At step 902, the
17 query optimizer 114 is invoked by the query processor 112 to determine if an
18 aggregation operation – a partial pre-aggregation – should be performed as a part
19 of the query. If it is determined that a partial pre-aggregation would not be
20 efficient ("No" branch, step 902), the several records are input to the record store
21 116. When the record store is full, a join operator is applied to the records in the
22 record store 116 at step 906. When the join has concluded processing, the records
23 in the record store 116 are no longer needed and the record store is cleared at step
24 908. If there are more records to process at step 210, then the process returns to
25 step 904 and another batch of input records are loaded into the record store 116

1 ("Yes" branch, step 910). If there are no more records ("No" branch, step 910),
2 then the process is terminated.

3 If it is determined that a partial pre-aggregation should be performed as part
4 of the query ("Yes" branch, step 902), then a single input record is received at step
5 912. The records in the record store 116 are searched to at step 914 to determine if
6 there is a record in the record store 116 having the same grouping column value as
7 the input record. In the present example, the grouping column is the customer
8 number field from the sales table. If there is a match between the input record and
9 a record in the record store 116 ("Yes" branch, step 916), then the input record is
10 combined with the matching record at step 916 and a new input record is received
11 at step 912. Combining records in the present example consists of adding the
12 invoice amount from the input record to the aggregate invoice amount of the
13 matching record in the record store. (Note that when an input record is combined
14 with a record in the record store, the number of records in the record store does not
15 increase and, therefore, does not affect the capacity of the record store.)

16 If there is no record in the record store 116 that matches the input record
17 ("No" branch, step 914), then at step 918, the input record is added to the record
18 store. At step 920, the record store 116 is checked to determine if it has reached
19 its capacity. If not ("No" branch, step 920), then the process reverts to step 912
20 and another input record is received.

21 If the record store 116 is full to capacity ("Yes" branch, step 920), then at
22 step 922, a join operation is performed on the records in the record store to
23 combine other data with the records. In the present example, the customer number
24 and aggregate invoice amount for each record in the record store is joined with the
25 customer contact and contact phone information from the customer table. The join

1 operation is any type of known join operation, such as a nested loop join, a simple
2 hash join, a classical hash join, etc. When the join operation is completed, the
3 records in the record store no longer needed and the record store is cleared. In the
4 present example, the groups of records output from the join are aggregated to
5 produce one output record for each customer number.

6 At step 926, it is determined if there are more records to input. If there are
7 more records to process ("Yes" branch, step 926) then the process reverts to step
8 912 and another input record is received for processing. If there are no more
9 records to process, then the process terminates.

10 Fig. 10 shows the records that result from aggregating the groups of records
11 output from the partial pre-aggregation/join procedure. Record 1 980 includes
12 customer number 986 ("9810"), customer contact 988 ("Adams"), contact phone
13 990 ("123-4567") and total sales 992 ("900"). Record 2 982 includes customer
14 number 986 ("9815"), customer contact 988 ("Bethard"), contact phone 990
15 ("345-6789") and total sales 992 ("200"). Record 3 984 includes customer
16 number 986 ("9821"), customer contact 988 ("Danner"), contact phone 990 ("987-
17 6543") and total sales 992 ("1000").

18 Note that the final records output from the query are identical in Fig. 8 and
19 in Fig. 12. But the implementation that utilized a combination of a pre-
20 aggregation and a join required less memory and processing to complete than the
21 first described implementation. The memory savings are on the order of one-half.

1 **Conclusion**

2 The implementations described herein provide a relational database
3 computer system with the means to optimize certain queries beyond what is
4 currently known in the art. The implementation of the partial pre-aggregation
5 technique is relatively simple, and the additional provision of the query optimizer
6 ensures that partial pre-aggregation will only be performed when it increases the
7 overall efficiency of the query.

8 Although the description above uses language that is specific to structural
9 features and/or methodological acts, it is to be understood that the invention
10 defined in the appended claims is not limited to the specific features or acts
11 described. Rather, the specific features and acts are disclosed as exemplary forms
12 of implementing the invention.
13
14
15
16
17
18
19
20
21
22
23
24
25

1 **CLAIMS**

2
3 **1.** A method for processing a database query according to at least one
4 grouping column value, the method comprising:

5 partially pre-aggregating records in a database to provide a result that
6 contains at least two records having like grouping column values; and

7 aggregating records derived from the partial pre-aggregation to provide a
8 result that contains records having unique grouping column values.

9
10 **2.** The method as recited in claim 1, wherein the partially pre-
11 aggregating further comprises:

12 maintaining a record store in memory, the record store having one record
13 for each different grouping column value encountered in the operation;

14 receiving a new record;

15 combining the new record with a record having the same grouping column
16 value, if such a record exists; and

17 adding the new record to the record store in the memory if there is no
18 record in the record store that has the same grouping column value as the new
19 record.

20
21 **3.** The method as recited in claim 2, further comprising:

22 adding additional new records to the record store until the record store
23 reaches a capacity such that it can accept no new records; and

24 outputting one or more records from the record store to a subsequent
25 database operator.

1
2 4. The method as recited in claim 3, wherein after the one or more
3 records have been output to the subsequent database operator, the adding and
4 outputting are repeated until there are no new records to process.

5
6 5. The method as recited in claim 4, wherein any records remaining in
7 the record store after there are no new records to process are output to the
8 subsequent database operator.

9
10 6. The method as recited in claim 3, wherein the subsequent database
11 operator is a join.

12
13 7. The method as recited in claim 1, further comprising estimating the
14 costs and benefits of the partial pre-aggregation, and partially pre-aggregating the
15 records only if the estimating indicates that the benefits are greater than the costs.

16
17 8. The method as recited in claim 1, wherein the partially pre-
18 aggregating includes utilizing a hashing function.

19
20 9. The method as recited in claim 1, wherein the partial pre-aggregating
21 creates a record store in memory, and wherein the method further comprises
22 utilizing the record store in memory for one or more other database operators.

1 10. One or more computer-readable media having computer-executable
2 instructions that, when executed by a computer, perform the method recited in
3 claim 1.

4
5 11. A computer programmed to perform the method recited in claim 1.

6
7 12. A relational database system, comprising:
8 memory for storing a record store, the memory having a portion available
9 for query processing;

10 a query processor coupled to the memory to process a query on the record
11 store according to at least one grouping column value, the query processor being
12 configured to partially pre-aggregate the record store to provide a result that
13 contains at least two data records that have like grouping column values; and

14 the query processor being further configured to aggregate data records
15 resulting from the partial pre-aggregation to provide an aggregation result that
16 contains data records, no two of the data records having a same grouping column
17 value.

1
2 **13.** The relational database system as recited in claim 12, wherein the
3 query processor being configured to partially pre-aggregate the record store
4 comprises the query processor being configured to:

5 maintain a record store in the volatile memory, the record store having one
6 record for each different grouping column value encountered in the partial pre-
7 aggregation;

8 receive an input record from the non-volatile memory;

9 combine the input record with a record in the record store that has the same
10 grouping column value, if there is such a record; and

11 adding the input record to the record store if there is no record in the record
12 store that has the same grouping column value as the input record.

13
14 **14.** The relational database system as recited in claim 13, wherein the
15 query processor is further configured to:

16 add additional input records to the record store in memory until the record
17 store reaches a capacity such that it can accept no more input records; and

18 output the records in the record store to a subsequent database operator.

19
20 **15.** The relational database system as recited in claim 14, wherein the
21 query processor is configured to:

22 continually add input records to the record store; and

23 output one or more records from the record store to a subsequent database
24 operator when the record store reaches a capacity such that it can accept no more
25 new records, or whenever there are no new records to process.

1
2 **16.** The relational database system as recited in claim 12, wherein the
3 query processor is further configured to perform a join on the records resulting
4 from the partial pre-aggregation.
5

6 **17.** The relational database system as recited in claim 12, wherein the
7 query processor is further configured to create a record store as a result of the
8 partial pre-aggregation and utilize the record store in processing of another
9 database operator.
10

11 **18.** The relational database system as recited in claim 12, further
12 comprising:
13

14 a query optimizer configured to estimate the costs and benefits of the query
15 processor performing a partial pre-aggregation; and
16

17 wherein the query processor performs the partial pre-aggregation only if the
18 query optimizer indicates that the benefits of the partial pre-aggregation are
19 greater than the costs of the partial pre-aggregation.
20

21 **19.** The relational database system as recited in claim 12, wherein the
22 query processor is further configured to utilize hashing to perform the partial pre-
23 aggregation.
24
25

1 20. The relational database system as recited in claim 12, wherein the
2 query processor is further configured to utilize hashing and partitioning to perform
3 the partial pre-aggregation.

4
5 ~~21.~~ A relational database computer program embodied on a computer-
6 readable medium, comprising:

7 partial pre-aggregation code to partially pre-aggregate data records
8 according to at least one grouping column value to provide a partial pre-
9 aggregation result having two or more records having like grouping column
10 values; and

11 aggregation code to aggregate data records in the partial pre-aggregation
12 result to provide an aggregation result having records with unique grouping
13 column values.
14
15
16
17
18
19
20
21
22
23
24
25

1
2 **22.** The relational database computer program as recited in claim 21,
3 wherein the partial pre-aggregation code is designed to:

4 maintain a record store in memory, the record store having one record for
5 each different grouping column value encountered in the partial pre-aggregation;

6 receive a new record;

7 combine the new record with a record in the record store having the same
8 grouping column value, if such a record exists;

9 add the new record to the record store no record in the record store that has
10 the same grouping column value as the new record;

11 continuously add additional new records to the record store until the record
12 store has reached a record store capacity;

13 output one or more records from the record store to a subsequent database
14 operator when the record store has reached the record store capacity; and

15 output the records in the record store when there are no new records to
16 process.

17
18 **23.** The relational database computer program as recited in claim 22,
19 further comprising database operator code that utilizes the record store for input.
20
21
22
23
24
25

1
2 **24.** A relational database computer program stored on a computer-
3 readable medium, the relational database computer program comprising computer-
4 executable instructions that, when executed on a computer, perform the following
5 steps:

6 receiving a stream of input records;
7 aggregating each input record in the stream as it is received to create a
8 record store;
9 joining records in the record store with other data; and
10 aggregating the records output from the join.

11
12 **25.** The relational database computer program as recited in claim 24,
13 wherein:

14 the record store has a capacity that is less than the number of records in the
15 stream of input records; and
16 the aggregating each input record is performed until the record store
17 reaches capacity.

18
19 **26.** The relational database computer program as recited in claim 24,
20 further comprising computer-executable instructions that, when executed by a
21 computer, perform the following steps:

22 determining if it is optimal to aggregate the input records prior to
23 performing the join; and
24 performing the aggregation prior to the join only if a determination is made
25 that it is optimal to perform an aggregation prior to the join.

1
2 **ABSTRACT**

3 A partial pre-aggregation database operation improves processing
4 efficiency of database queries by reducing the number of records input into a
5 subsequent database operation, provided the query includes a final aggregation. A
6 query optimizer is provided to determine when it is economical to partially pre-
7 aggregate data records and when it is not. The partial pre-aggregation creates a
8 record store in memory as input records are received. The record store is then
9 used by another database operator, which saves the other database operator from
10 having to re-create the record store.

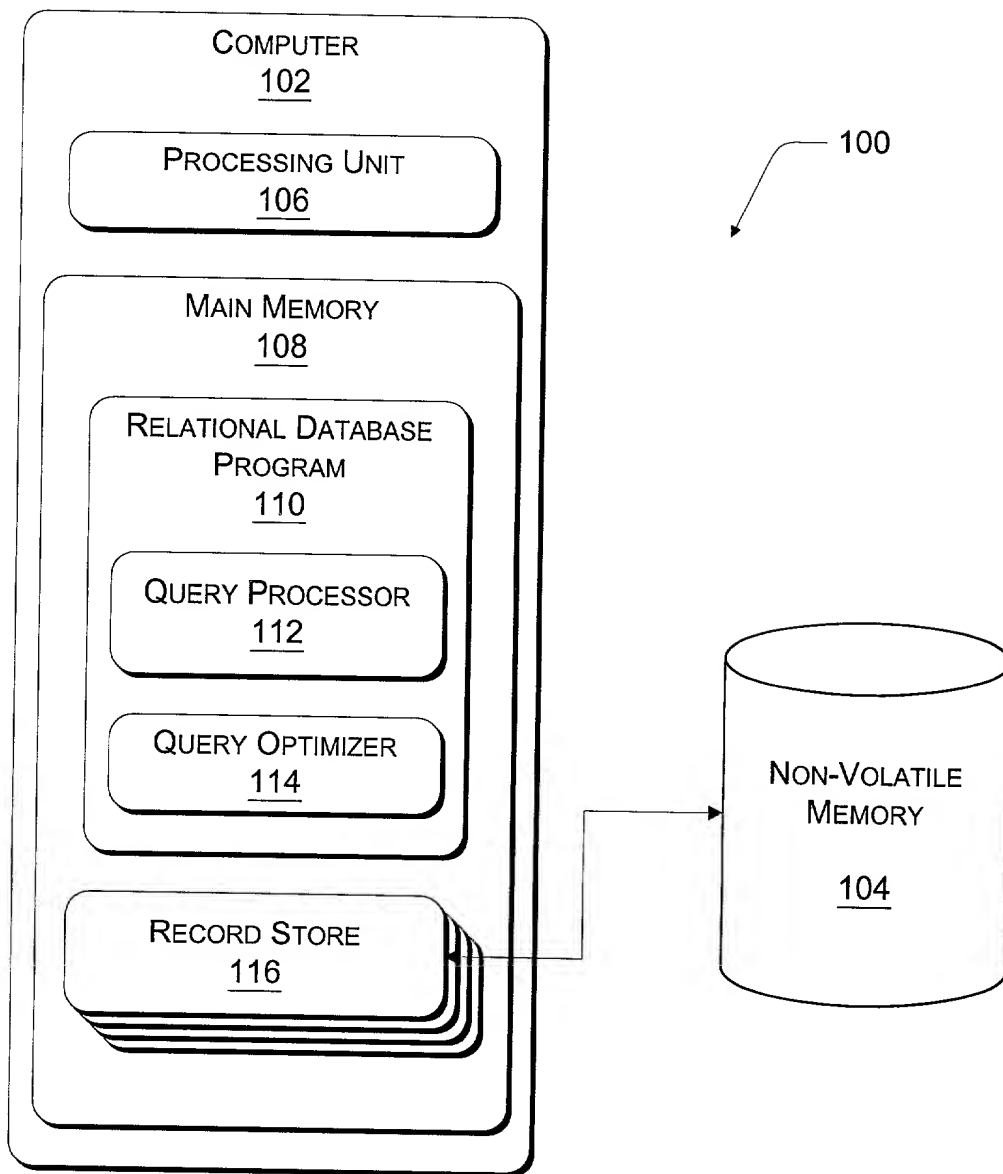


Fig. 1

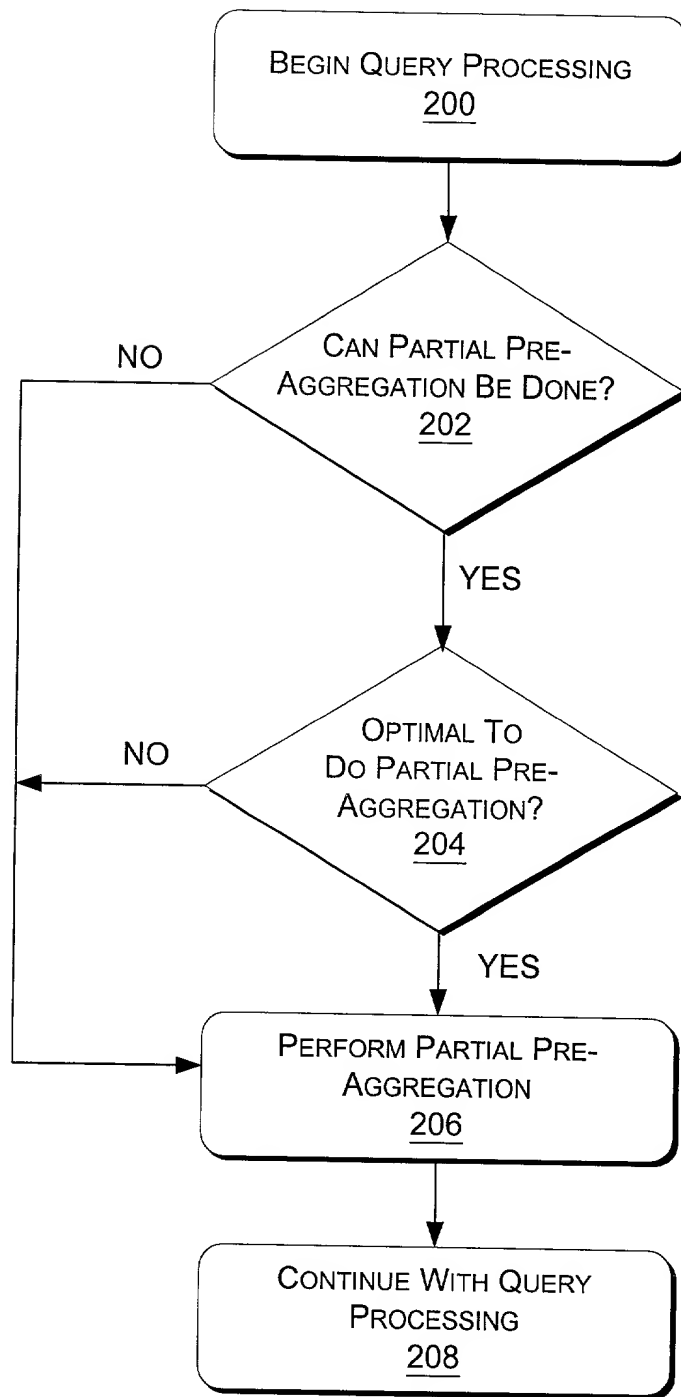


Fig. 2

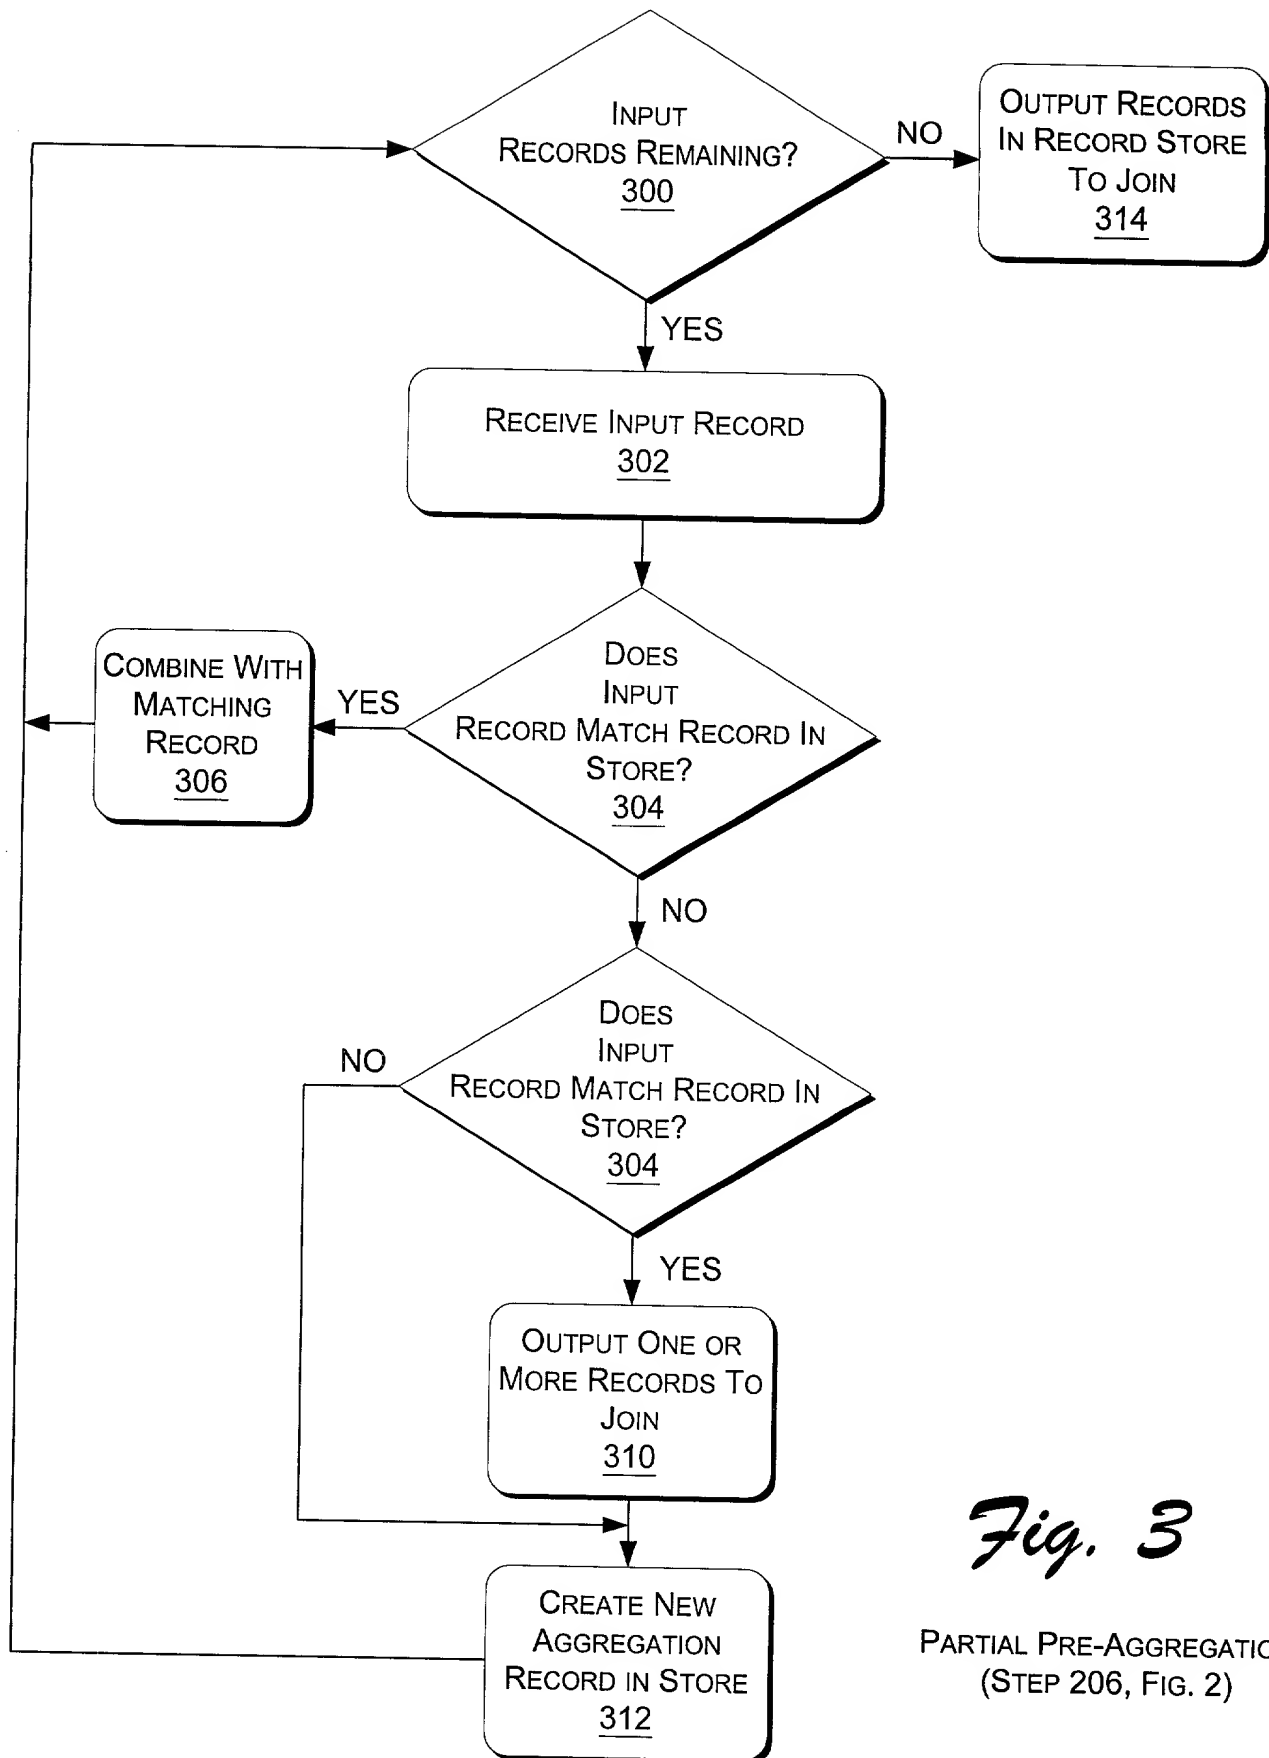


Fig. 3

PARTIAL PRE-AGGREGATION
(STEP 206, FIG. 2)

| | Customer Number | Invoice Date | Due Date | Invoice Amount |
|-------------|--------------------|-----------------|----------|-------------------|
| Record 1 | 9810 | | | \$100 |
| Record 2 | 9815 | | | \$200 |
| Record 3 | 9810 | | | \$300 |
| Record 4 | 9821 | | | \$400 |
| Record 5 | 9810 | | | \$500 |
| Record 6 | 9821 | | | \$600 |

Sales Table
(Customer Invoices)

Fig. 4

| | Customer Number | Customer Name | Customer Contact | Contact Phone |
|-------------|--------------------|------------------|---------------------|------------------|
| Record 1 | 9810 | | Adams | 123-4567 |
| Record 2 | 9815 | | Bethard | 345-6789 |
| Record 3 | 9819 | | Gill | 987-6543 |
| Record 4 | 9821 | | Danner | 765-4321 |

Customer Table
(Customer Contacts/Country)

Fig. 5

| | Customer Number | Total Invoice Amount | |
|-------------|--------------------|----------------------------|-----|
| Record 1 | 9810 | \$400 | 602 |
| Record 2 | 9815 | \$200 | 604 |
| | 606 | 608 | |

First Group of Aggregation Records
Output To Join Operator

Fig. 6a

| | Customer Number | Total Invoice Amount | |
|-------------|--------------------|----------------------------|-----|
| Record 1 | 9821 | \$1000 | 610 |
| Record 2 | 9810 | \$500 | 612 |
| | 614 | 616 | |

Second Group of Aggregation Records
Output To Join Operator

Partial Pre-Aggregation Results

Fig. 6b

| | Customer Number | Customer Contact | Contact Phone | Total Sales |
|-------------|--------------------|---------------------|------------------|----------------|
| Record 1 | 9810 | Adams | 123-4567 | \$400 |
| Record 2 | 9815 | Bethard | 345-6789 | \$200 |
| Record 3 | 9821 | Danner | 987-6543 | \$1000 |
| Record 4 | 9810 | Adams | 123-4567 | \$500 |

Join Results

700

Fig. 7

| | Customer Number | Customer Contact | Contact Phone | Total Sales |
|----------|-----------------|------------------|---------------|-------------|
| Record 1 | 9810 | Adams | 123-4567 | \$900 |
| Record 2 | 9815 | Bethard | 345-6789 | \$200 |
| Record 3 | 9821 | Danner | 987-6543 | \$1000 |

Aggregation Results

800

| | Customer Number | Customer Contact | Contact Phone | Total Sales | |
|-------------|--------------------|---------------------|------------------|----------------|-----|
| Record 4 | 9810 | Adams | 123-4567 | \$900 | 980 |
| Record 5 | 9815 | Bethard | 345-6789 | \$200 | 982 |
| Record 6 | 9821 | Danner | 765-4321 | \$1000 | 984 |
| | 986 | 988 | 990 | 992 | |

Result Of Aggregation On Records
Output From Record Store

Fig. 10

1 **IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

2 Inventorship.....Larson et al.
 3 Applicant Microsoft Corporation
 4 Attorney's Docket No. MSI-479US
 5 Title: Partial Pre-Aggregation In Relational Database Queries

6 **DECLARATION FOR PATENT APPLICATION**

7 As a below named inventor, I hereby declare that:

8 My residence, post office address and citizenship are as stated below next to
 9 my name.

10 I believe I am the original, first and sole inventor (if only one name is listed
 11 below) or an original, first and joint inventor (if plural names are listed below) of the
 12 subject matter which is claimed and for which a patent is sought on the invention
 13 entitled "Partial Pre-Aggregation In Relational Database Queries," the specification
 14 of which is attached hereto.

15 I have reviewed and understand the content of the above-identified
 16 specification, including the claims.

17 I acknowledge the duty to disclose information which is material to the
 18 examination of this application in accordance with Title 37, Code of Federal
 19 Regulations, § 1.56(a).

20 PRIOR FOREIGN APPLICATIONS: no applications for foreign patents or
 21 inventor's certificates have been filed prior to the date of execution of this
 22 declaration.

23 **Power of Attorney**

24 I appoint the following attorneys to prosecute this application and transact all
 25 future business in the Patent and Trademark Office connected with this application:
 26 Lewis C. Lee, Reg. No. 34,656; Daniel L. Hayes, Reg. No. 34,618; Allan T.

1 Sponseller, Reg. 38,318; Steven R. Sponseller, Reg. No. 39,384; James R.
2 Banowsky, Reg. No. 37,773; Lance R. Sadler, Reg. No. 38,605; Michael A. Proksch,
3 Reg. No. 43,021; Thomas A. Jolly, Reg. No. 39,241; David A. Morasch, Reg. No.
4 42,905; Kasey C. Christie, Reg. No. 40,559; Katie E. Sako, Reg. No. 32,628 and
5 Daniel D. Crouse, Reg. No. 32,022.

6 Send correspondence to: LEE & HAYES, PLLC, 421 W. Riverside Avenue,
7 Suite 500, Spokane, Washington, 99201. Direct telephone calls to: James R.
8 Banowsky (509) 324-9256.

9
10 All statements made herein of my own knowledge are true and that all
11 statements made on information and belief are believed to be true; and further that
12 these statements were made with the knowledge that willful false statements and the
13 like so made are punishable by fine or imprisonment, or both, under Section 1001 of
14 Title 18 of the United States Code and that such willful false statement may
15 jeopardize the validity of the application or any patent issued therefrom.

16
17 * * * * *

18 Full name of inventor: Per-Ake Larson

19 Inventor's Signature  Date: 6/28/2000

20 Residence: Redmond, WA


21 Citizenship: Canadian

22 Post Office Address: 14923 NE 76th Court
23 Redmond, WA 98052
24
25

Full name of inventor:

Cesar A. Galindo-Legaria

Inventor's Signature



Date: 6/28/2000

Residence:

Redmond, WA

Citizenship:

Mexico

Post Office Address:

16141 NE 44th Court
Redmond, WA 98052